

REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-06-0059

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 28-02-2006		2. REPORT TYPE FINAL		3. DATES COVERED (From - To) FROM 15-02-02 TO 15-12-05	
4. TITLE AND SUBTITLE Assessing and Improving Team Decision Making				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER AFOSR F49620-02-1-0119	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER	
6. AUTHOR(S) Fischler, Ira; Sorkin, Robert D.				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Dr. Ira Fischler, Dept. Psychology, P.O. Box 112250, Gainesville, FL 32611-2250				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Dr. Jerome Busemeyer, Cognitive Program, Air Force Office of Scientific Research AFOSR/NL, 875 North Randolph Street, Suite 325 Arlington, Virginia 22203-1768				10. SPONSOR/MONITOR'S ACRONYM(S) AFOSR/NL	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approve for Public Release: Distribution Unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This project used analytical and experimental techniques derived from signal detection theory to quantify the decision making performance of individuals and teams. The basic decision task was to decide on the presence or absence of signals in noise. The project's experiments studied how individual and team performance depends on member signal-to-noise ratio, correlation among member inputs, efficiency of member updating of likelihood estimates, and constraints on member interaction and communication. The results show that: (a) members combine individual estimates with high efficiency to form the team's decision, relative to an optimal Bayesian rule, (b) team performance depends on the response protocol that constrains the sequential order of information exchange among team members, and (c) team members choose to acquire additional sources of information (varying in signal-to-noise ratio, bias, cost, time, and correlation) in a near optimal manner.					
15. SUBJECT TERMS signal detection, team consensus, team collaboration, selection of information sources					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 18	19a. NAME OF RESPONSIBLE PERSON Ira Fischler
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code) (352) 392-0601 x228

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Assessing and Improving Team Decision Making

FINAL PERFORMANCE REPORT

AFOSR/NL
Grant F49620-02-1-0119

February 15, 2002 - December 14, 2005

20060309 075

February 15, 2006

Ira Fischler
University of Florida
Gainesville FL 32611-2250

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

1. TECHNICAL SUMMARY

ABSTRACT

This project used analytical and experimental techniques derived from signal detection theory to: (1) quantify the decision-making performance of individuals and teams; particularly with respect to how information received from sources having different properties is chosen, weighed and combined, (2) model the decision behavior of experienced teams of human operators operating under incentives for decision speed and accuracy, and (3) describe how the accuracy and speed of group deliberation depends on the aggregation rule and response protocol that constrains the sequential order of information exchange among team members. The basic (individual or group) decision task was to decide on the presence or absence of signals in noise. Signals were presented to operators on individual graphical displays and the team then had to reach a decision about signal occurrence. The project's experiments show how individual and team performance depends on team member signal-to-noise ratio, correlation among members' inputs, efficiency of member updating of likelihood estimates, and constraints on member interaction and communication. Specifically, the project experiments demonstrated that: (a) members' individual estimates are combined to form the team's decision with high efficiency, relative to the optimal Bayesian rule, (b) team performance depends, in a predictable way, on the response protocols that constrain the sequential order of information exchange among team members, and (c) team members choose (and evaluate) additional sources of information that vary in signal-to-noise ratio, bias, cost, time, and correlation in a near optimal fashion. These results have important potential applications for optimizing team decision making, particularly for (a) structuring member communication so as to minimize the effects of interruptions, biases, correlations, and other constraints, (b) enabling the optimal aggregation of information for team decisions at maximum accuracy and minimal cost and time, (c) quantifying and training expert team behavior, and (d) designing the most effective systems for spatially separated, networked teams of human operators. These results were reported in two Master's theses, two doctoral dissertations, and a number of research articles and presentations.

Basic Theory and Experimental Procedure

This research (see, Sorkin, Luan & Itkowitz, 2004) employed an approach derived from electrical engineering analyses of distributed (or decentralized) signal detection (see Pete et al., 1993a,b; Sorkin & Dai, 1994; and Viswanathan & Varshney, 1997) to model individual and team decision behavior. This approach has generally not received attention in the traditional human judgment and decision literature.

In the classic signal detection situation, the decision maker (DM) makes an observation, x , and then must decide whether her observation was caused by a signal-plus-noise or a noise-alone event (hereafter referred to simply as 'signal' and 'noise'). Her ability to discriminate between the two events is defined by the two distributions, $f(x|n)$, the likelihood that the observation was due to noise, and $f(x|s)$, the likelihood that the observation originated from signal. An ideal (or optimal) decision rule is the likelihood ratio, $l(x) = f(x|s)/f(x|n)$. DM's ability to discriminate between the two events is given by the detection parameter, d' , the normalized separation between the means of the two (assumed) normal distributions, as shown in figure 1. All of the information that DM has about the existence of a signal on any trial is specified by the value of the likelihood ratio, $l(x)$. These distributions are sometimes plotted on a log-likelihood ratio axis; this preserves their normal shape, generalizes the dimensions of the decision axis, and simplifies subsequent updating of the likelihood estimate by additional observations.

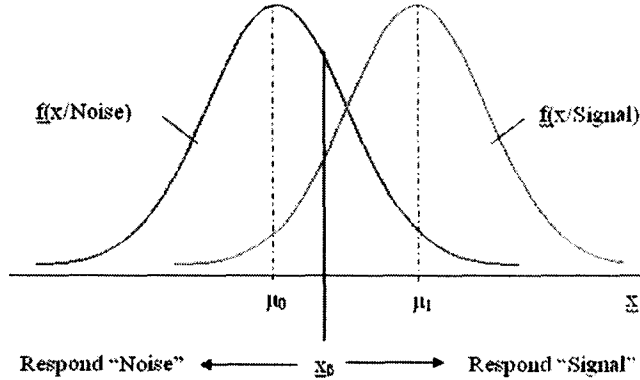


Figure 1.

By defining the payoffs (or penalties) for the four possible decision outcomes on a trial (correct identification of signal or "hit", correct identification of noise or "correct rejection", incorrect signal response or "false alarm", incorrect noise response or "miss"), one can define an optimal criterion (or threshold) for the value of $\ln l(x)$. That optimal rule is

$$\text{respond "signal" iff } \ln l(x) \geq \ln[\beta] \quad [1]$$

where β is defined by (Tanner & Swets, 1954)

$$\beta = \frac{V_{\text{correct-rejection}} + V_{\text{false-alarm}}}{V_{\text{hit}} + V_{\text{miss}}} \cdot \frac{p(\text{noise})}{p(\text{signal})} \quad [2]$$

where the V_{ij} are the utilities of the four possible outcomes, and $p(\text{noise})/p(\text{signal})$ is the prior odds ratio of noise to signal.

The Advice Task

Suppose that additional information were available from an uncorrelated, external advisor, A, a source with detection sensitivity d'_A and criterion β_A . (For discussion of correlated sources and information aggregation, see Durlach, Braida, & Ito, 1986; Sorkin & Dai, 1994). How and when should information from advisor A be incorporated into the DM's decision? In the continuous advice case, the advisor provides an estimate of signal occurrence that is summarized by its likelihood estimate, $l_A(x)$. All of the information present is captured by combining DM's likelihood estimate with the advisor's likelihood estimate, in $\ln[l(x_{DM}, x_A)]$. If the DM and A sources are independent, the joint likelihood estimate is given by the product of the likelihood estimates and

$$\ln[l(x_{DM}, x_A)] = \ln[l_{DM}(x) \cdot l_A(x)] = \ln[l_{DM}(x)] + \ln[l_A(x)] \quad [3]$$

$$\text{Then DM's decision rule is respond "signal" iff } \ln[l_A(x)] \geq \ln \beta - \ln[l_{DM}(x)] \quad [4]$$

The resulting hit probability is given by the volume under the bivariate normal *signal* distribution, $f\{\ln[l_{DM}(x), \ln[l_A(x)]]|s\}$, defined by the area above the diagonal criterion line (a line with slope equal to -1 and crossing the abscissa at $\ln(\beta)$). The false alarm probability is given by the (similar) volume under the noise distribution, $f\{\ln[l_{DM}(x), \ln[l_A(x)]]|n\}$.

If accessing the information has a cost, DM establishes advice purchase criteria on $\ln[l_{DM}(x)]$. That is, DM only purchases the advice when her observation is above a value $\ln \beta_1$ and below a value $\ln \beta_2$. The optimal settings for these criteria can be determined by calculation of the expected value of the trial outcome minus the advice cost.

A number of interesting cases result when the information from A is binary rather than continuous (advisor A's response is either "signal" or "noise"). One such case is the so-called "confirmation bias" problem: DM is faced by the option of choosing one of two advisors: advisor

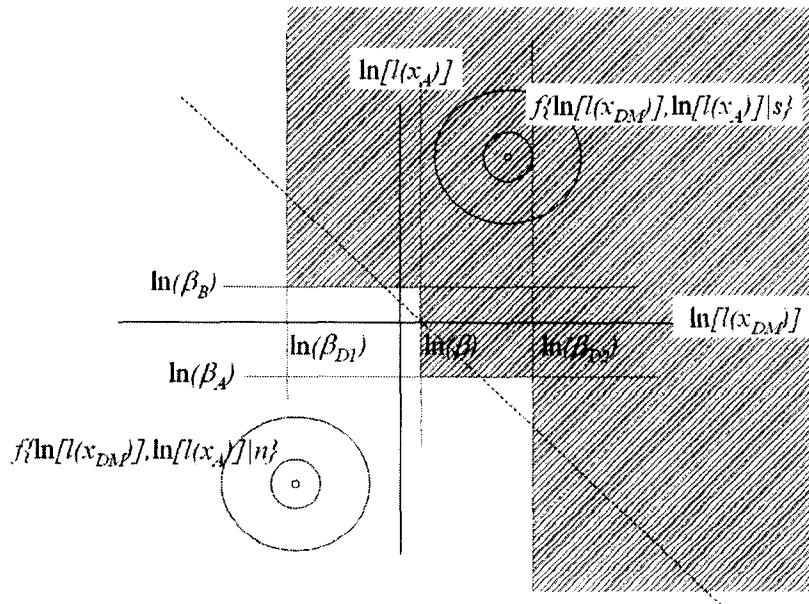


Figure 2. Binary advice situation.

A or B. Both advisors have the same cost and expertise; the only difference is the response criteria that they employ, $\ln(\beta_A)$ and $\ln(\beta_B)$. Advisor A uses a liberal (low) criterion, and B uses a conservative (high) criterion. The question is how DM's choice of advisor should depend on the magnitude of DM's initial estimate. It can be shown that the optimal choice after making a high estimate is to choose the liberal advisor, and the optimal choice after having made a low estimate is to choose the conservative advisor. This is the so-called "confirmation" strategy, wherein the advisor's recommendation is *usually* expected to support DM's initial estimate.

The situation can be understood graphically by examining figure 2. As in the continuous case, DM's goal is to maximize the volume of the correct regions under the respective signal and noise distributions. How should the DM's choice of an advisor depend on the location of her likelihood estimate when her estimate is greater than $\ln(\beta_{D1})$ but less than $\ln(\beta_{D2})$? The respective response regions for the two advisors are indicated by the horizontal lines in Figure 2; the liberal advisor responds "signal" when his estimate is above $\ln(\beta_A)$, and the conservative advisor responds "signal" when his estimate is above $\ln(\beta_B)$. The following consulting strategy is optimal: If the DM's estimate is between $\ln(\beta_{D1})$ and $\ln(\beta)$, she should consult the conservative advisor $\ln(\beta_B)$, if her estimate is between $\ln(\beta)$ and $\ln(\beta_{D2})$, she should consult the liberal advisor $\ln(\beta_A)$. It can be seen that the alternative or disconfirming strategy, of reversing the choice of liberal and conservative advisors, would result in a lower probability of correct responses. (Note the optimum criterion line overlaid on the diagram. The square area with sides $[\ln(\beta_A) \text{ to } \ln(\beta_B)]$ by $[\ln(\beta) \text{ to } \ln(\beta_{D2})]$ would be un-shaded in the disconfirming strategy, and the area with sides $[\ln(\beta_A) \text{ to } \ln(\beta_B)]$ by $[\ln(\beta_{D1}) \text{ to } \ln(\beta_D)]$ would now be a shaded region resulting in an increase in false alarm rate and decrease in hit rate.) The optimal strategy can be interpreted as a "confirmation" strategy, since on most trials the DM will choose the advisor who is most likely to agree with her initial estimate. An intuitive explanation of this is to realize that if DM has made a high evaluation, only a negative (and unlikely) reply from the liberal advisor should cause her to reject her initial estimate.

Team Decision Tasks

Our approach to modeling group decision making shares many features with the advice or information acquisition problem. We assume that group performance depends not only on the members' individual expertise, but also on how member information is communicated, aggregated, and converted to the group decision. Members may have different levels of expertise (effective signal-to-noise ratios), biases (response criteria), correlations (degrees of shared noise), and response rules. We assume that members generally update their estimates in an optimal Bayesian fashion. However, the communication of member information is assumed to be inherently imperfect. Figure 3 shows a general diagram of such a system. The team of local decision makers (LDMs) must decide about the occurrence of a signal; the task is made difficult by noise (both common and shared) and by the LDMs' having a discrete response vocabulary. After receiving the input, each member estimates the likelihood of signal occurrence and converts that estimate to a categorical response r_{ij} . That response is communicated to the Decision Center (DC) where the group's final decision is made. The DC uses a set algorithm, such as a specific majority of signal or noise votes, to aggregate the members' information. The DC also may employ a deliberation process that involves additional interaction and feedback among the members.

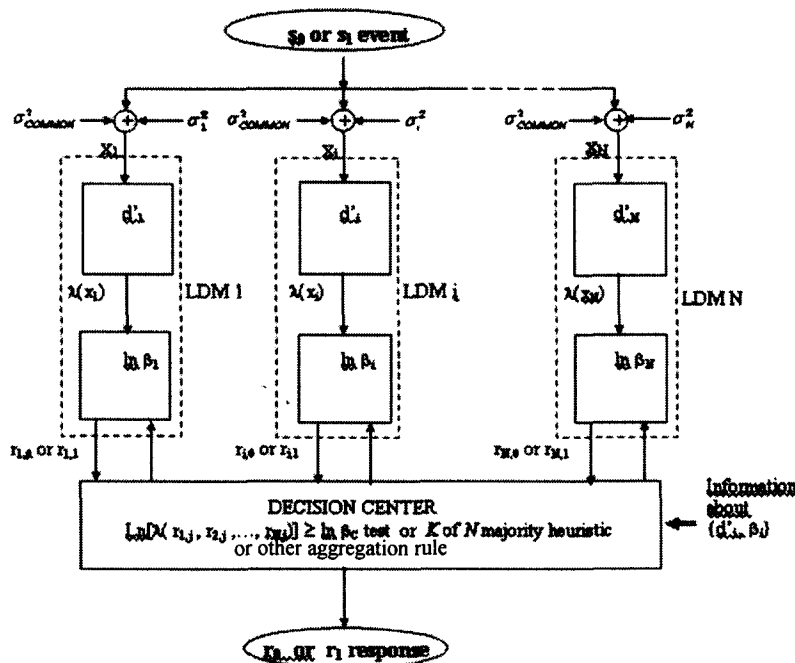


Figure 3. [$\lambda(x_i)$ indicates likelihood ratio.]

In a hypothetical jury, each member updates her opinion (and re-votes) upon hearing the simultaneous, non-anonymous votes of the other jurors. All members vote at once, member judgments are updated in an optimal fashion, and the updating and voting process is repeated until a majority or stopping criterion is reached. The updating calculation of the juror's response is equivalent to the juror's Bayesian recalculation of likelihood ratio. It is assumed that each juror is knowledgeable about the expertise and bias of the other jurors. Under some conditions, jury performance can closely approximate the maximum likelihood aggregation of the members' continuous responses (Swaszek & Willett, 1995; Sorkin et al., 2004).

An interesting and important deviation from the jury aggregation model occurs when the LDM information arrives sequentially rather than simultaneously; such as in a group of electronically networked LDMs. We've tested the effects of different sequential response rules on simulated group performance and have shown that the order of arrival of member information has a potentially large effect on the accuracy and speed of the decision. We assume a member's vote is conveyed to the group in an individual, sequential fashion. As each member's response is made, the other members update their individual estimates in a Bayesian fashion. A fixed aggregation rule is then used to either arrive at a group decision via a hidden vote. Members don't speak again until all have spoken. This 'deliberation' process continues until a group decision is made or time runs out. Table 1 shows some possible protocols that could be used to determine the next team member to speak.

Table 1.

The next speaker to access the network is:

Random	$Ran(i)$
Most/least expert	$Max(d'), Min(d')$
Most/least biased	$Max \ln \beta_i , Min \ln \beta_i $
Largest/smallest estimate	$Max \ln \lambda(x_j) , Min \ln \lambda(x_j) $
Most/least confident	$Max \ln \lambda(x_j) - \ln \beta_i , Min \ln \lambda(x_j) - \ln \beta_i $
Largest/smallest weight	$Max w_i(r_j)$ where: $w_i(r_d) = \ln[1 - p(r_{i,1} s_d)] / [1 - p(r_{i,1} s_d)]$ $w_i(r_d) = \ln[p(r_{i,1} s_d)] / [p(r_{i,1} s_d)]$

Experimental Method

The basic decision task is to judge whether the stimulus present on an experimental trial was generated by a signal-plus-noise or noise-alone condition. Each participant is presented with a graphical display like the set of analog gauges shown in figure 4, and the group must report whether the displays were due to a signal or noise. The setting of each display (vertical) element is determined by a normal distribution whose mean depends on the parameters of the experiment. Versions of this task have been studied extensively in our laboratory (Luan, Sorkin, & Itzkowitz, 2004; Montgomery and Sorkin, 1996; Sorkin & Dai, 1994; Sorkin, Hays, & West, 2001; Sorkin, Mabry, Weldon, & Elvers, 1991; and Sorkin, West, & Robinson, 1998). The difficulty of the task is determined by the display's physical and statistical parameters. The major physical parameters are the display duration and the visual angle subtended by the display. The statistical factors are the means of the signal-plus-noise and noise-alone distributions and the value of their common standard deviation. In our experiments, the main factor that determines task difficulty is the display signal-to-noise ratio (SNR), which is the difference between the distribution means divided by their (common) standard deviation. The individual SNR for each member can be controlled, as can the correlation between member displays. We record the mean value, x_{ij} , of the trial-by-trial display presented to each member i on each trial j . The monetary payoff to the participants depends on the accuracy of the *group's* detection performance. Sometimes the monetary payoff also depends on the accuracy of the individual members' performance on their first response; this manipulation serves to control their initial response criterion. This general decision task can be run under different rules for member interaction and information display. More complex displays can be employed to convey the expertise and criterion of information sources and the correlation of different sources.

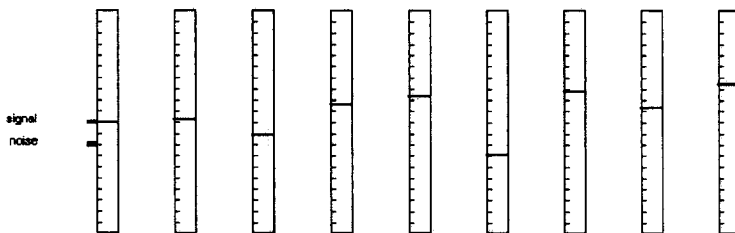


Figure 4. Example of a typical display (signal-plus-noise trial).

Results

Effect of Consistency and Correlation of Information Sources

This experiment tested how decision makers weigh the estimates received from different sources when those sources vary in their (a) number, (b) reliability or expertise, (c) apparent consistency (internal correlation), and (d) correlation with the decision maker's initial estimate. The human decision makers performed the graphical decision task, aided by simulated team members. The decision maker observed a display of a signal-plus-noise or noise-alone event and then made an estimate of the likelihood of signal occurrence. The decision maker was then supplied with a graphical display of several estimates made by one or two sets of simulated team members. The decision maker was then required to make a final yes-no decision about the occurrence of signal on that trial. The payoff to the decision maker was based on the accuracy of the final yes-no decision. Performance was assessed when the estimates of the virtual team members were either pair-wise uncorrelated or correlated (among themselves) at some level, or had differing overall indices of detection. By computing the correlation over trials between the decision and the source's magnitude, the experiment assessed how much decision weight the decision maker gave to each information source.

One experimental condition tested which of two equal-information sources (sources with equal aggregate- d') would be given the higher weight: the one with the internal pair-wise correlation and higher individual d 's, or the one with the zero pair-wise correlation and lower component d 's. Thus, this condition tested for the presence of a bias toward sub-source consistency or sub-source expertise. Further conditions were run comparing sources with differing overall informational value. The results of these experiments indicated that there is a small but significant bias toward information sources that have higher consistency and higher component expertise, even though the information available from such sources is identical to *or less* than that received from lower consistency sources. This is an important result because it shows the existence of a here-to-fore unreported bias toward information sources that have higher component expertise or internal consistency. Such a bias will result in inefficiency in decision performance; performance will decrease as a function of the discrepancy between actual and optimal weighting of input sources. It may be possible to provide a decision maker with the estimates received from different sources so that the overall information display adjusts for such a pre-existing weighting bias.

Effect of Correlation Between Decision-maker and External Information Sources

When information from two external sources is available to help a decision-maker (DM) make her final decision, following the majority choice of the two sources and decision-maker can be a useful decision heuristic for rapidly achieving a final decision. The advantage of this heuristic is greatest when the information from the three sources arises from independent observations and the sources are equal in expertise. However, when one of the source's information is correlated with the DM's (which would result in more agreement between that source and the DM), the statistical advantage of information aggregation is reduced. Therefore, as the correlation between one source and the DM increases, the DM should assign less final decisional weight to that source (and herself), relative to the estimate from the other, independent, source. In this experiment, two information sources with nearly equal detection abilities to the participants were presented after the DM made an initial estimate. We manipulated the correlation (ρ_{PA}) between one information source (A) and the participants' decision estimates at three different levels: 0, 0.4 and 0.7, and kept the other source (B) independent ($\rho_{PB} = 0$,

$\rho_{AB} = 0$). The results indicated that when ρ_{PA} increased, participants (1) assigned significantly more decision weight to information from themselves and source A, and failed to assign appropriate decision weight to source B; (2) continued to use the simple majority decision rule to make their final decision and usually did not follow source B's estimates when B's estimates were counter to the estimates of the participant and source A; and (3) consequently had reduced final decision efficiencies as ρ_{PA} was increased. We concluded that the simple majority decision heuristic was extensively used when decision-makers integrated information from different sources and that they weighed information from agreeing (and correlated) sources much more than information from disagreeing sources. This type of bias results in a loss in the benefit of information contained in the disagreeing source as well as in a reduced accuracy in the final decision.

Information Acquisition

Seeking advice from other people or sources is a common practice in making real-life decisions and in command and control decision making. We developed a normative model that describes how advice-taking should depend on the decision maker's own estimate, as well as the cost, expertise, and decision bias of the potential advisors. Three experiments were conducted to examine how human participants took and utilized advice in different decision environments, and whether their observed behaviors were consistent with the model's prescriptions. In the experiments, a decision maker (DM) first observed a visual stimulus and then made an initial estimate of the nature of the stimulus: signal-plus-noise (S) or noise-alone (N). Then, information from external sources was available for the DM to purchase regarding the decision event. After the DM made her consulting decision and integrated the advice with her own, she made a final decision. On each trial of one experiment, participants were asked to make a decision about whether or not to consult an advisor. Both the cost (High or Low) and the displayed content of the advice were manipulated (Continuous or Binary). It was found that participants consulted the advisor more frequently when their own estimates were uncertain and also when the advice was expressed in the continuous mode rather than the binary mode.

In a further experiment, two advisors with equal expertise and consulting costs but different decision biases (criteria, β) were presented. Four hypothesized consulting strategies were first tested in computer simulations: (a) a Low Estimate followed by a choice of a Liberal Advisor or a High Estimate followed by a choice of a Conservative Advisor, (b) a Low Estimate followed by a choice of a Conservative Advisor or a High Estimate followed by a choice of a Liberal Advisor, (c) Always choose Liberal Advisor, and (d) Always choose Conservative Advisor. The (a) strategy of "Low-C-High -L", which has a similar working rationale as the so-called "confirmation-bias" strategy, was the best in simulations and theory (see figures 2.,5. and earlier theoretical description of the Advice problem). This strategy prescribes that: if a DM has an observation that is more likely to lead her to make a N decision, she should consult a "Conservative" advisor; otherwise, she should consult a "Liberal" one. Because a Conservative advisor was defined in this study as the advisor who was biased to make more N decisions than S (tended to give more N than S opinions to the DM) and a Liberal advisor was defined to do the contrary, "Low-C-High-L" is actually a "confirming" strategy: a DM should always consult the advisor who is more likely to agree with her own decision estimate. The results are shown in figure 6. Most subjects followed the optimal strategy; none followed the non-optimal strategy. It can be concluded that people have certain behavioral tendencies when taking advice from external sources and that these generally confirm to optimal behavior. Our advice model can be used to identify precisely how those observed behaviors deviate from the optimal, and may contribute additional insight into understanding human advice-acquisition and decision-making.

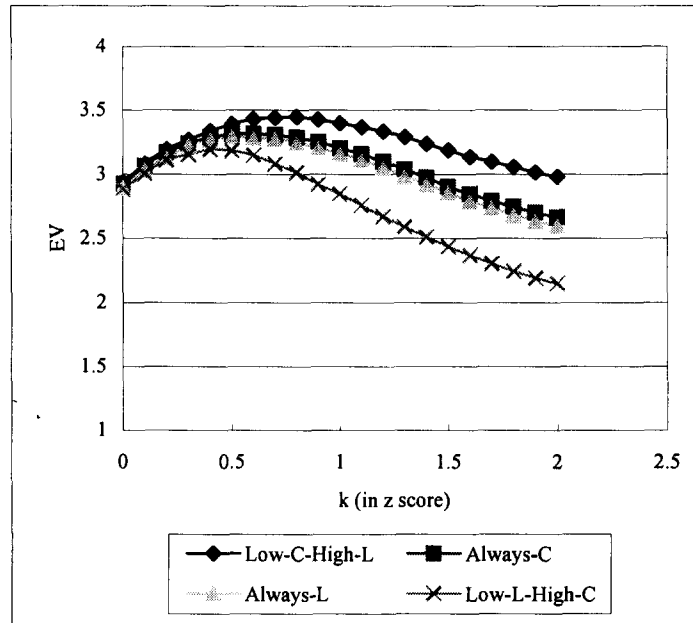


Figure 5. The expected payoff (EV) as the function of k (an index of consulting frequency) resulted from the four tested strategies in the optional consulting condition. The optimal strategy has the highest payoff; always choosing one source is in between; the reverse strategy is the worst.

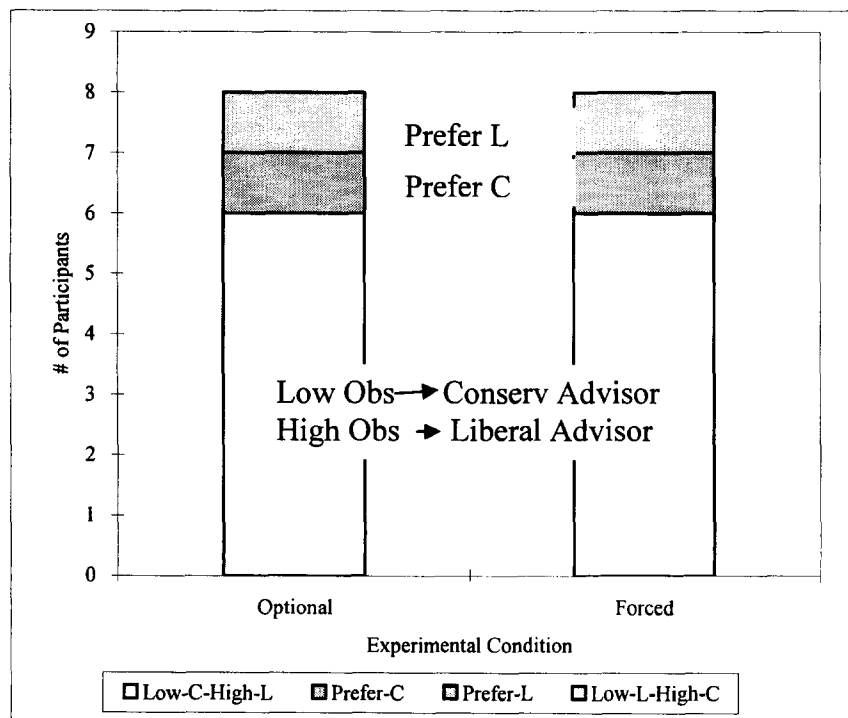


Figure 6. The number of participants who used each consulting strategy in the two experimental conditions of the advisor experiment. (None of the participants used the Low-L-High-C strategy).

Team Deliberation

A series of experiments studied collaborative detection in decision making groups. Collaboration is defined as the process by which group members sequentially share their private information with the group. The many factors that influence the group during this process can be grouped into two categories: group composition (the accuracies, biases, and correlation of members' opinions) and organizational structure (speaking rules, response protocols, and decision criteria). Using a normative model, we studied the effects of the team decision criteria (simple majority and unanimity) and the response protocols (member properties determining who speaks next), as well as the effects of group composition and organizational structure on performance. Real and simulated group performance was evaluated in a number of ways (accuracy, post-deliberative bias, deliberation length, and deliberation efficacy). The model makes some very specific assumptions: members attempt to optimize the group payoff/performance, members are rational decision makers with appropriate motivation, and members update their estimates and combine information from other group members' public votes in a Bayesian manner. The general architecture of the model is shown in figure 7.

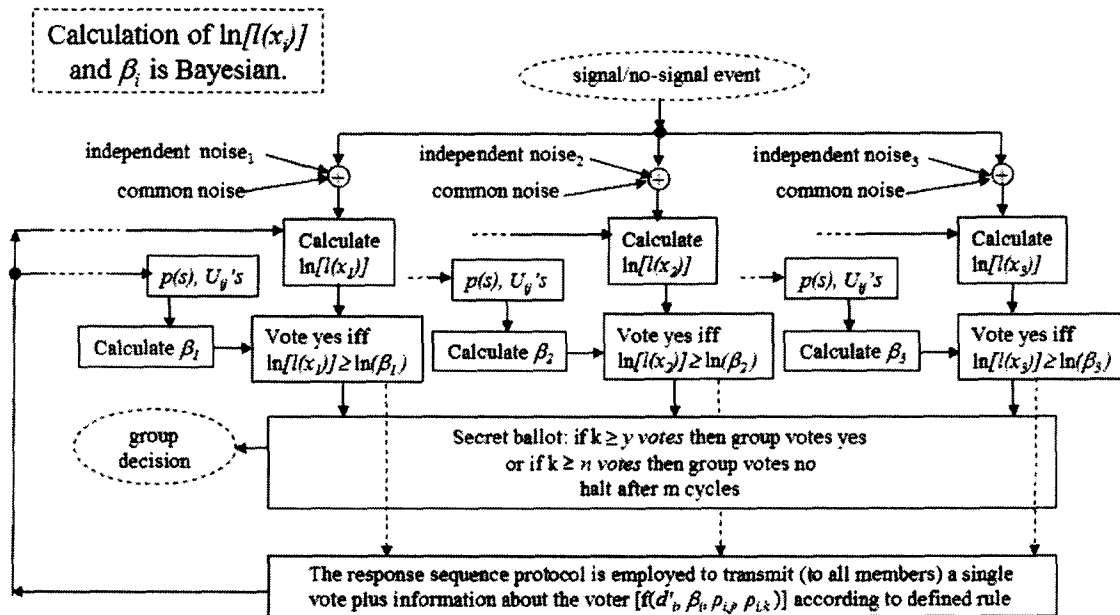


Figure 7. Computational architecture of group decision process.

One experiment tried to determine which, if any, response protocols were used in groups when members participated according to a voluntary structure. Because we desired speaking orders that reflected group members' desire to share their information with the group, we set the governing response protocol to sequence speakers in ascending order of their response time to their initial stimulus. Previous research indicated that group members who are certain (or confident) of their decision are the most likely to participate (Carlston, 1977; Sneizek & Henry, 1989). We found that a members' observation magnitude, relative to their own individual decision criterion, had a small but clear impact on the speed of a group members' decision.

Team deliberation, even with a requirement for unanimity, was relatively rapid. This was consistent with the predictions of Swaszek & Willett (1995). Second, group performance was better than that of the best member in the group, providing evidence consistent with the model

(and simulations) and contrary to some previous human studies that found that groups rarely perform better than their best member (Gigone & Hastie 1997; Einhorn, Hogarth, & Klemperer, 1977). Third, group deliberation with a requirement for unanimity had the effect of “debiasing” the most extreme group members. Even when two-thirds of the group members began deliberation with an initial bias (either liberal or conservative), the group decision had a generally neutral bias.

All of the groups tested had higher decision performance than that of the group’s best member, indicating that the groups were using the information provided in deliberation. None of the protocols produced results having *both* the highest detection accuracy and the minimum deliberation duration. Thus, these two performance metrics are not optimized by the same protocol. Simulated and actual groups using a Minimum d' rule and a unanimous decision criterion had higher detection ability than the other groups. Simulated and actual groups (see figure 8) requiring unanimity and using a Minimum d' protocol, had the slowest performance, e.g., highest number of votes than other groups.

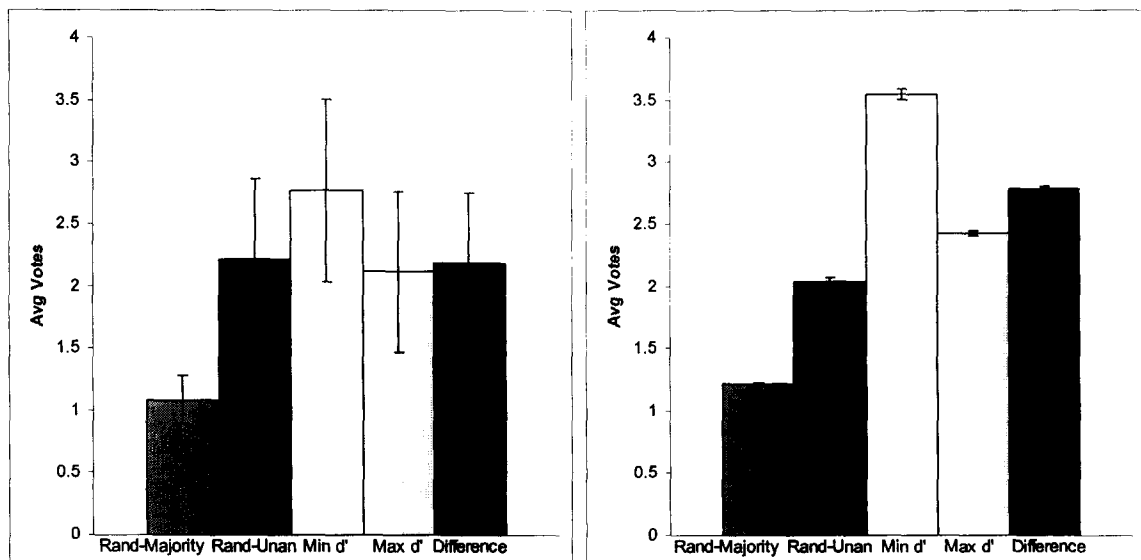


Figure 8. Mean deliberation times and associated standard errors for actual groups (left panel) and simulated groups (right panel). The left-most column is for a simple majority rule; the right four columns are for unanimous rule groups using, respectively, a random, minimum d' , maximum d' and maximum observation-criterion protocol.

Conclusions and Possible Applications

It is clear that team performance is a function both of the team members’ characteristics and the organizational structure used to reach consensus decisions. Whereas the former mostly affects individuals’ private judgments, the latter affects how these judgments are integrated with information received from other team members. Our approach provides an integrated model for addressing individual and team decision making behavior in deliberative groups and provides a consistent account of such behavior in the laboratory. In general, individuals within the groups tested updated their own opinions according to the normative, Bayesian, predictions. Moreover, groups closely matched the predictions of group accuracy and deliberation length provided by simulated groups. Although our experimental procedure was simplified and did not reflect the complexity of actual deliberative teams, our results indicate that groups act in a largely rational way.

The basic issues addressed by our model (and the additional questions raised here) have important implications. Because so many important decisions are made by both face-to-face and networked deliberative groups (e.g., juries, military teams, and corporate boards), it is essential that we discover ways to maximize the performance of these groups. In some situations, it is clear that adding some automated control of the “deliberation” process in a networked group could result in small but measurable increases in accuracy or important decreases in decision time.

A. PERSONNEL

Ira Fischler (interim P.I.)
Professor, Department of Psychology
University of Florida

Jesse Itzkowitz
Graduate Student, Department of Psychology
University of Florida

Shenghua Luan
Post Doctoral Assignment
Max Planck Institute for Human Development
Berlin, Germany

Robert D. Sorkin (P.I.)
Professor, Department of Psychology
University of Florida

B. PUBLICATIONS, THESES, & DISSERTATIONS

Adamowicz, W., Hanemann, M., Swait, J., Johnson, R., Layton, D., Regenwetter, M., Reimer, T., & Sorkin, R. Decision Strategy and Structure in Households: A “Groups” Perspective. *Marketing Letters*, 2005, **16**, Issue 3-4, 387-399.

Luan, S. (2002). Information utilization from outside sources in decision-making. M.S. Thesis, University of Florida.

Luan, S. *To consult or not: A signal detection study of advice-taking*. Ph.D. Dissertation University of Florida, 2004.

Luan, S., Sorkin, R.D., and Itzkowitz, J. Weighing information from outside sources: A biased process. *Journal of Behavioral Decision Making*, 2004, **17**, 95-116.

Luan, S., Sorkin, R.D., & Itzkowitz, J. A unified model of advice acquisition and use. *Psychonomic Bulletin and Review* (submitted).

Sorkin, R.D., Luan, S., and Itzkowitz, J. Group decision and deliberation: A distributed detection process, Chapter 23 in Nigel Harvey & Derek Koehler (Eds.), *Handbook of Judgment and Decision Making*. Oxford: Blackwell, 2004.

Itzkowitz, J. *Bias in individuals and groups revisited: Using Signal Detection Theory to examine the effects of individual member bias on group bias and group accuracy*. M.S. thesis, University of Florida, 2003.

PUBLICATIONS, THESES, & DISSERTATIONS (continued):

Itzkowitz, J. *Improving the performance of deliberative groups through changes in organizational structure: A signal detection approach*. Ph.D. Dissertation, University of Florida, 2005.

Sorkin, R. D., Itzkowitz, J., Luan, S., & Crandall, C. *The American jury: An optimal system for group decision-making*. (in preparation)

Itzkowitz, I, Sorkin, R. D., and Luan, S. Effects of response order on the performance of deliberative groups. (in preparation)

C. PAPERS AT PROFESSIONAL MEETINGS

Itzkowitz, J., Sorkin, R. D., & Luan, S. *Effect of speaking order on group deliberation*. Annual Meeting of the Society for Judgment and Decision Making, Vancouver, Canada, November 2003.

Luan, S., Sorkin, R. D., & Itzkowitz, J. *Is There a Decision Bias For Information From Internally Consistent Sources?* Annual Meeting of the Cognitive Science Society, Fairfax Virginia, August, 2002.

Luan, S., Sorkin, R. D., & Itzkowitz, J. *Signal Detection Model of Advice Acquisition*, 45nd Annual Meeting of the Psychonomic Society, Minneapolis, MI, November 2004.

Luan, S., Sorkin, R. D., & Itzkowitz, J. *The majority heuristic in information integration*. Annual Meeting of the Society for Judgment and Decision Making, Vancouver, Canada, November 2003.

Sorkin, R.D., Luan, S., & Itzkowitz, J. *A rational model of advice acquisition*. Annual Meeting of the Society for Advancement in Behavioral Economics, Philadelphia, PA., 2004.

Sorkin, R. D., Luan, S., & Itzkowitz, J. *Effect of Majority Rule and Initial Bias on Information Aggregation by Groups*. AF, NSF, and EADM Workshop on Information Aggregation, Silver Spring, MD, May 2003.

Sorkin, R. D. *ARGU: Architectures for Rational Group Use*. AFOSR and Rensselaer Polytechnic Institute Workshop on Integrated Models of Cognitive Systems (IMCS), Saratoga Springs, NY, March, 2005.

Sorkin, R. D., *Group Decisions: Analyzing Decision Strategy and Structure in Households and Firms*. 6th Annual CU-Boulder Invitational Choice Symposium, June, 2004.

D. INTERACTIONS & CONSULTATIONS (R. Sorkin)

Member, Editorial Board of *Human Factors & Ergonomics*

Member, Executive Board, American Psychological Association Division 21 (Engineering Psychology), and Chair, Division 21 Fellows Committee.

Talks: Sorkin, R. D. *Modeling Group Decision Making*, Dept. Elect. & Comp. Eng., University of Connecticut, Storrs, Ct., November 2003; and Dept. Psychology, University of Texas, Austin, TX, February 2004.

Workshop on Detecting Deception in Language and Cultural Context, Center for the Advanced Study of Language (CASL), University of Maryland, College Park, MD, 14-25 June, 2004.

E. DISCOVERIES/INVENTIONS None

F. HONORS/AWARDS

Appointment of Robert Sorkin as Senior Scientist for Cognitive Science & Engineering, Air Force Research Laboratory, WPAFB, Ohio, November 2005.

REFERENCES

- Birnbaum, M. H. & Mellers, B. A. (1983). Bayesian inference: combining base rates with opinions of sources who vary in credibility. *Journal of Personality and Social Psychology*, 45, 792-804.
- Carlston, D. E. (1977). Effects of polling order on social influence in decision-making groups. *Sociometry*, 40(2), 115-123.
- Clemen, R. T. (1989). Combining Forecasts: A review and annotated bibliography. *International Journal of Forecasting*, 5, 559-583.
- Connolly, T. (1988). *Studies of information-purchase processes. Chapter 12 in Human Judgment: The SJT view / edited by Brehmer B. & Joyce C.*, Elsevier Science Pub. Co.
- Durlach, N. I., Braida, L. D. & Ito, Y. (1986). Towards a model for discrimination of broadband signals. *Journal of the Acoustical Society of America*, 80, 63-72.
- Edwards, W. (1965). Optimal strategies for seeking information. *Journal of Math. Psych.*, 2, 312-329.
- Einhorn, H. J., Hogarth, R. M., & Klempner, E. (1977). Quality of group judgment. *Psychological Bulletin*, 84, 158-172.
- Ferrell, W. R. (1985). Combining individual judgments. In G. Wright (Ed.), *Behavioral Decision Making* (pp. 111-145). New York: Plenum.
- Gigone, D., & Hastie, R. (1997). Proper analysis of the accuracy of group judgments. *Psychological Bulletin*, 121, 149-167.
- Grofman, B., & Feld, S. (1986). Determining optimal weights for expert judgment. In B. Grofman & G. Owen (Eds.) *Decision Research* (Vol. 2, pp 167-172), Greenwich, CT: JAI Press.
- Guarnaschelli, S., McKelvey, R.D., & Palfrey, T.R. (2000). An experimental study of group decision rules. *American Political Science Review*, 94, 407-423.
- Gustafson, D. H., Shukla, R. K., Delbecq, A., & Walster, G. W. (1973). A comparative study of differences in subjective likelihood estimates made by individuals, interacting groups, Delphi groups, and nominal groups. *Organizational Behavior and Human Performance*, 9, 280-291.
- Hastie, R. (1986). Experimental evidence on group accuracy. In B. Grofman & G. Owen (Eds.), *Decision Research* (Vol. 2). Greenwich, CT: JAI Press.

- Hey, J. D. (1981). Are Optimal Search Rules Reasonable? And Vice Versa? (And Does It Matter Anyway?). *Journal of Economic Behaviour and Organization*, 2, 47-70.
- Hoballah, I. Y., & Varshney, P. K. (1989). Distributed Bayesian signal detection. *IEEE Transactions on Information Theory*, 35, 5, 995-1000.
- Hung, A. A. & Plott, C. R. (2001). Information cascades: Replication and an extension to majority rule and conformity-rewarding institutions. *The American Economic Review*, 91, 1508-1520.
- Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94, 211-228.
- Larson, J. R., Jr., Sargis, E. G., & Bauman, C. W. (2004). Shared knowledge and subgroup influence during decision-making discussions. *Journal of Behavioral Decision Making*, 17, 245-262.
- Luan, S., Sorkin, R. D. & Itzkowitz, J. (2004). Weighing information from outside sources: A biased process. *Journal of Behavioral Decision Making*, 17, 95-116.
- Montgomery, D. A. & Sorkin, R. D. (1996). Observer sensitivity to element reliability in a multi-element visual display. *Human Factors*, 38, 484-494.
- Nitzan, S., & Paroush, J. (1982). Optimum decision rules in uncertain dichotomous choice situations. *International Economic Review*, 23, (2) 289-297.
- Penner, L. A. & Davis, J. (1969). Conformity and the "rational" use of unanimous majorities. *Journal of Social Psychology*, 78, 299-300.
- Penrod, S., & Hastie, R. (1979). Models of jury decision making: A critical review. *Psychological Bulletin*, 86, 3, 462-492.
- Pete, A., Pattipati, K.R., & Kleinman, D. L. (1993a). Optimal team and individual decision rules in uncertain dichotomous situations. *Public Choice*, 75, 205-230.
- Pete, A., Pattipati, K.R., & Kleinman, D. L. (1993b). Distributed detection in teams with partial information: A normative-descriptive model. *IEEE Transactions on Systems, Man, and Cybernetics*, 23, 1626-1647.
- Shapley, L., & Grofman, B. (1984). Optimizing group judgmental accuracy in the presence of interdependencies. *Public Choice*, 43, 329-342.
- Snizek, J. A. (1992). Groups under uncertainty: An examination of confidence in group decision making. *Organizational Behavior and Human Decision Processes*, 52, 124-155.
- Snizek, J. A., & Buckley, T. (1995). Cueing and cognitive conflict in judge-advisor decision making. *Organizational Behavior and Human Decision Processes*, 62, 159-174.
- Snizek, J. A., & Henry, R. A. (1989). Accuracy and confidence in group judgment. *Organizational Behavior and Human Decision Processes*, 43, 1-28.
- Sorkin, R. D. & Dai, H. (1994). Signal detection analysis of the ideal group. *Organizational Behavior and Human Decision Processes*, 60, 1-13.
- Sorkin, R. D., Hays, C. J., & West, R., (2001). Signal-detection analysis of group decision making. *Psychological Review*, 108, 183-203.
- Sorkin, R. D., Luan, S. & Itzkowitz, J. (2004). Group decision and deliberation: A distributed detection process. In Koehler, D., & Harvey, N. (Ed.), *Blackwell Handbook of Judgment and Decision Making*. (pp. 464-484). London: Blackwell.
- Sorkin, R. D., Mabry, T. R., Weldon, M., & Elvers, G. (1991). Integration of information from multiple element displays. *Organizational Behavior and Human Decision Processes*, 49, 167-187.
- Sorkin, R. D., West, R. and Robinson, D. E. Group performance depends on the majority rule. *Psychological Science*, 1998, 9, 456-463.
- Sorkin, R., & Dai, H. (1994). Signal detection analysis of the ideal group. *Organizational Behavior and Human Decision Processes*, 60, 1-13.
- Sorkin, R., Hays, C., & West, R. (2001). Signal-detection analysis of group decision making. *Psychological Review*, 108, 183-203.

- Stasser, G., Kerr, N. L., & Davis, J. H. (1989). Influence processes and consensus models in decision-making groups. In P. Paulus (Ed.), *Psychology of group influence* (2nd Ed.). Hillsdale, NJ: Erlbaum.
- Stasser, G., Vaughan, S. I., & Stewart, D. D. (2000). Pooling unshared information: The benefits of knowing how access to information is distributed among members. *Organizational Behavior and Human Decision Processes*, 82, 102-116.
- Swaszek, P. E., & Willett, P. (1995). Parley as an approach to distributed detection. *IEEE Transactions on Aerospace and Electronic Systems*, 31, 1, 447-457.
- Tanner, W. P., Jr., & Swets, J. A. (1954). A decision-making theory of visual detection. *Psychological Review*, 61, 401-409.
- Viswanathan, R. & Varshney, P. K. (1997). Distributed detection with multiple sensors: Part I—Fundamentals. *Proceedings of the IEEE*, 85, 54-63.
- Yaniv, I. (2004). Receiving other peoples' advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, 93, 1-13.

PERFORMANCE REPORTS

All agreements require Performance Reports. Performance Reports may be yearly or final reports. Performance Reports will be due no later than 1 September each year that the award is active. However, agreements effective after 31 July will not require a Performance Report until 1 September of the following year. Final Performance Reports are due 90 days after the expiration of the agreement. All Performance Reports will be submitted to the Program Manager. Performance Reports are described below:

PERFORMANCE REPORT

1. Cover Sheet: As a minimum, the cover sheet should include the following information: Principal Investigator's name, Institution's name and address, and agreement number.

2. Objectives: List the objectives of the research effort or the statement of work: This may be omitted if there has been no change. State new or revised objectives if they have changed and the reason why.

3. Status of effort: A brief statement of progress towards achieving the research objectives. (Limit to 200 words).

4. Accomplishments/New Findings: Describe research highlights, their significance to the field, their relationship to the original goals, their relevance to the AF's mission, and their potential applications to AF and civilian technology challenges.

5. Personnel Supported: List professional personnel (Faculty, Post-Docs, Graduate Students, etc.) supported by and/or associated with the research effort.

6. Publications: List peer-reviewed publications submitted and/or accepted during the 12-month period starting the previous 1 October (or since start for new awards).

7. Interactions/Transitions:

a. Participation/presentations at meetings, conferences, seminars, etc.

b. Consultative and advisory functions to other laboratories and agencies, especially Air Force and other DoD laboratories. Provide factual information about the subject matter, institutions, locations, dates, and name(s) of principal individuals involved.

c. Transitions. Describe cases where knowledge resulting from your effort is used, or will be used, in a technology application. Transitions can be to entities in the DoD, other federal agencies, or industry. Briefly list the enabling research, the laboratory or company, and an individual in that organization who made use of your research.

8. New discoveries, inventions, or patent disclosures. (If none, report None.)

9. Honors/Awards: List honors and awards received during the grant/contract period. List lifetime achievement honors such as Nobel Prize, honorary doctorates, and society fellowships prior to this effort.

10. Markings: In order to ensure prompt receipt and acceptance, mark the outside of the package clearly to indicate that it is a performance report.

FINAL PERFORMANCE REPORT

The purpose of the final Performance Report is to document and to transition the results of the effort into the Air Force and DoD applied research community. The final report will always be sent to the Defense Technical Information Center (DTIC) and unclassified reports will be available to the public through the National Technical Information Service (NTIS).

Content: The final report is more than an extension of previous progress reports. **The final report shall be a comprehensive technical summary of the significant work accomplished.** The final report, where it is not readily accessible in published form, should where applicable: 1) Clearly describe and illustrate the experimental equipment, set up, and procedures; 2) Characterize and tabulate collected/computed data in an appendix; and 3) Sufficiently describe computational codes so they can be reproduced. Include a listing of the code in an appendix if possible and appropriate.

When the research effort culminates in the production of one or more student theses or dissertations, in these cases, the most significant advancements and conclusions (equations, figures, relationships, etc.) should be included in an executive summary. The theses or dissertations should be attached as appendices only if they are not readily available. If they are, clearly reference them and how they can be obtained. Also include in the executive summary, cumulative lists of people involved in, and publications stemming from, the research effort. Do not include copies of already submitted or published articles in the final report.

Format: Cover and title page. Standard Form (SF) 298, Report Documentation Page, shall be used. Item 13 of the form should contain a 100 to 200 word abstract summarizing technical progress during the reporting period. Style should be third person singular using past tense. Jargon, special symbols or notations, subscripts, mathematical symbols or foreign alphabet letters are not permitted. The pages should be prepared for acquisition and distribution by DTIC. All pages should be of good quality for copying purposes. No pages should be missing.

Grantees are encouraged to submit reports that are printed/copied double sided on recycled paper that has at least 20% postconsumer material. The format and standard required by your institution for the preparation of theses and dissertations shall be used for the final report. In the absence of any institutional standards, you may wish to refer to the National Information Standards Organization (NISO) document Z39.18-1995, "Scientific and Technical Reports--Elements, Organization, and Design," for guidance. The publication may be obtained from the organization below:

NISO Press Fulfillment
P.O. Box 338
Oxon Hill, MD 20750-0338 USA
Phone: 301.567.9522 or 1.800.282.6476
Fax: 301.567.9553
Internet Address: <http://www.niso.org>

In order to ensure prompt receipt and acceptance, mark the outside of the package clearly to indicate that it is a final performance report.